

CSCI 315 Operating Systems Design

Midterm Exam 2 Study Guide

- Review labs assignments, in-class activities, and quizzes. Make sure that you have a solid understanding of the topics they address.
 - Review the points in the chapter reading guides linked to the class schedule.
 - This document doesn't mean to give an exhaustive coverage of what might appear in the exam, but it will be useful as a self-check list for your preparation.
-

1. Identify the following concepts.

- multiprogramming
- process
- thread
- context switch
- scheduling
- memory protection
- starvation
- mutual exclusion
- atomicity
- semaphore (counting semaphore; binary semaphore/mutex)
- deadlock
- resource allocation graph
- deadlock prevention, deadlock avoidance, deadlock recovery
- demand paging
- swapping
- thrashing
- first-fit, best-fit, worst-fit
- frame, page
- virtual memory
- internal fragmentation, external fragmentation
- physical address, logical address
- TLB
- page table, hierarchical page table

- Belady's anomaly
- memory mapped file
- valid bit, dirty bit
- allocation of frames (global vs. local, fixed vs. priority)
- effective access time

2. How can one estimate the length of CPU bursts so that this information can be used in scheduling of processes and/or threads?
3. Explain how each of the following scheduling algorithms works. Be able to discuss how each one affects CPU utilization, job throughput, turnaround time, waiting time, and response time. Discuss how preemption might apply to each of the algorithms and what effects it might produce on the performance of the system.
 - first come, first served
 - shortest job first
 - shortest remaining time first
 - round robin
 - priority
 - multilevel queue
4. How does the use of preemption and priority levels affect the behavior and the performance of the scheduling algorithms above?
5. Given a set of process with specified arrival times and CPU burst lengths, for one or more scheduling algorithms: draw a Gantt chart of the execution of the processes, calculate performance metrics of the scheduling algorithm, apply the performance metrics to reason on how well the scheduler works and propose alternative scheduling for the given scenario.
6. What are the four necessary conditions for a deadlock to occur?
7. Given a set of processes, resources, and resource requests, draw the resource allocation graph after each request and state whether or not a deadlock is possible.
8. In general, a cycle in a resource allocation graph indicates only the possibility of a deadlock. Under what special condition does it indicate the existence of a deadlock?

9. What are the different ways of dealing with deadlock?
10. How does the *safety* version of the Banker's algorithm work? You should be able to carry out a hand execution of the algorithm.
11. In contiguous memory allocation schemes, it is possible to use first-fit, best-fit, and worst-fit strategies. State a justification for each of these strategies. Is there one that works better or worse than the others? Explain your answer.
12. Identify the circumstances in which internal fragmentation and external fragmentation happen.
13. What would it take for a system with contiguous memory allocation to minimize external fragmentation? Would the solutions you propose have benefits that outweigh their implementation and operational costs?
14. What are the advantages and disadvantages of memory management schemes such as: overlays, swapping, and virtual memory? What is the impact that each of these schemes have on the usability of the system (from a programmer's perspective) and on the implementation of the system?
15. Propose a mechanism (hardware, software, or a combination of both) by which programs can be loaded anywhere in memory.
16. In the context of a paging system, what is a logical address? What is a physical address?
17. Describe what a *Translation Lookaside Buffer* (TLB), what data it contains, and what it does for a paging system.
18. Describe the impact that the use of a TLB can have on the effective access time of a virtual memory system.
19. Explain how a physical address is determined in a paged system using a TLB. Be specific about what the values are and how they are used.
20. Given the logical address in a virtual memory system, describe how it is translated to a physical address when there is: a single page table, a hierarchical page table.

21. Discuss the pros and cons of using a single page table and a hierarchy of page tables.
22. Consider a paged system where addresses are byte addresses, and pages consist of 16 4-byte words. If the desired page is in frame 5, and the offset is byte 4 in the frame, what is the corresponding physical address?
23. In a system with a TLB, assume the following:
 - a. the memory access time is 150 nsec.
 - b. the TLB access time is 25 nsec.
 - c. the TLB hit rate is 80%.Computing the effective access time for memory in this type of scenario.
24. What are the motivations for using virtual memory in an operating system?
25. Identify the benefits of using virtual memory in a multi-user, multi-programmed operating system.
26. Identify the steps the OS takes in handling a page fault.
27. Describe a mechanism by which one frame of physical memory can belong to the logical address space of multiple different processes.
28. Construct a scenario in which it is better to pre-load all the pages of a process than to allow pages to be loaded on demand.
29. Describe the advantages and disadvantages of fixed and priority frame allocation in a virtual memory system.
30. Describe the advantages and disadvantages of global and local page replacement in a virtual memory system.
31. Describe what causes *thrashing* and propose effective strategies to avoid or to stop it.

32. Given a sequence of logical address references, identify which references causes page hits and page faults according to the following page replacement algorithms: optimal, FIFO, LRU, MFU, and LFU.
33. Compare different ways to implement and to approximate LRU page replacement.
34. Describe the purpose and the operation of second hand page replacement, identifying what hardware support it may need.
35. Describe the concept of *working set* for processes executed in a virtual memory system.